



Predicting Staff Sizes to Maintain Networks

Dr. Lon D. Gowen
The MITRE Corporation

MITRE completed a three-month study to assess the state of the practice in staffing levels for maintaining a computer-networking infrastructure (CNI). The state of the practice was determined by looking at technical papers on the subject, conducting organizational and technical-expert surveys, and looking at software models that attempt to predict staffing levels. There were very few quantitative heuristics available in the literature; however, the data did show that typical CNIs have a 1:42 ratio of support staff to users. That is, one full-time equivalent of CNI staffing per 42 users for a typical CNI. This number can vary, up or down, by 17 percent or more depending on the details of the CNI. The Department of Defense, as well as the private sector, can use the results of this study to predict initial CNI support levels, to support their current level of staffing, or to justify an increase or decrease in staffing. Additionally, this paper breaks down CNI support into four major areas: systems administration, hardware maintenance, help desk, and configuration management, and provides ratios for predicting each of them within a typical CNI.

Having the appropriate manpower to maintain a given computer-networking infrastructure (CNI) is an important factor to consider since only 26 percent of a local access network's total cost of ownership (TCO¹) is hardware, while the remaining 74 percent is labor [1]. Of the 74 percent for labor, typically 43 percent is for end-user operations, 17 percent for technical support, and 14 percent for administration [1]. Other common reasons for having an accurate, up-to-date figure deal with budgeting, reliability, and quality. If an organization's CNI staffing levels are too high, then it wastes resources. If the staffing levels are too low, then response times, reliability, and end-product quality suffers; overtime is too high; and workers leave for a better working environment (since at present, the demand is considerably greater than the supply).

Every organization within the Department of Defense (DoD) has to

estimate staffing levels (manpower) for maintaining their CNI. This activity takes place on a regular basis for most

“Therefore, even if the data accurately portrays the state of the practice, it may not portray the optimum, since the state of the practice may not be optimal.”

organizations whether part of the DoD or the private sector. While not a fascinating or appealing topic of research for many people, manpower-sizing predictions are a critical part of planning. Therefore, at the request of its DoD sponsors, MITRE conducted a three-

month study to try to determine the current state of the practice in CNI staffing levels. The focus was primarily on the private sector. However, the study also looked at some DoD-based data.

This article is a sanitized version of the study's full report [2]. Due to the full report's sensitive nature and critical views in certain areas, this article maintains the anonymity of informational sources. Such an approach was necessary to obtain honest data.

For purposes of this study, MITRE determined the state of the practice by collecting information from the following sources: recent technical papers (most within 24 months), organizations currently supporting CNIs, technical experts currently working in the field, and current modeling tools.

Current CNI Staffing Practices

This section lists the data that MITRE collected from the four sources mentioned previously. After discussing the data from these four areas, a section removes the outliers to show how well the data tightens up. There are some people who will have problems with dropping the outliers for statistical reasons; however, the whole purpose of dropping the outliers is not to present any kind of statistical proof. Instead, it is merely to show the effect such changes have on the averages and standard deviations of the remaining data. Such comparisons are very useful to certain organizations.

Table 1 is a list of the data MITRE collected. The next few sections reference this data in more detail; however, there are a few things worth mentioning. First, as the table shows, the deviations are too

Table 1: CNI Staffing Data Collected

Type of Data	Source of Data	Number of Users Per FTE of CNI Support			
		Systems Administration	Help Desk	Hardware Maintenance	Configuration Management
Technical Papers	1. Lucent INS [3]	155.2	113.8	284.5	853.5
	2. Gartner #1 [4]	106.7	77.6	106.7	
	3. Gartner #2 [5-6]	247.8	60.0		
	4. PC Week [7]		86.0		
	5. IDC [8]		99.0		
Org. Surveys	6. DoD	103.3	110.7	106.9	442.9
	7. Private A	80.0	80.0	80.0	
	8. Private B	71.0		71.0	
Technical Expert Surveys	9. DoD Sector	426.8	81.3	213.4	284.5
	10. Private Sector	227.6	136.6	162.6	227.6
COTS Modeling Tools	11. Run A	92.1	376.0	305.9	388.0
	12. Run B	80.8	199.4	193.1	292.8
Mean (Average)		159.1	129.1	169.3	414.9
Standard Deviation		112.9	90.2	86.6	228.4
Percent Standard Deviation		71.0	69.9	51.1	55.1

loose. Dropping a few outliers within the four areas of CNI support can tighten them. The numbers, however, still show useful similarities. For example, the averages between systems administration, help desk, and hardware maintenance are relatively close to each other but considerably smaller than configuration management. Second, the statistical means for all four areas seem reasonable and in general agreement. Third, the standard deviations (as a percent of the mean) are very high. When combining data from the four areas (see Table 2), the percentage drops drastically. While MITRE did not investigate this drop, there are two potential reasons: an inadequate understanding of one or more of the four areas by some or all of the sources, and a different operational definition of these terms.

Each ratio in Table 1 represents how many users one full-time equivalent (FTE) of CNI staffing can support for a given area. For example, Lucent's paper [3] recommends one help-desk FTE per 113.8 users. The shorthand for such a ratio, in this article, is 1:113.8.

Technical Papers

Of the 29 papers reviewed, only six papers [3-8] contained sufficient information to be useful for predicting CNI support levels. MITRE collected the six papers into five groups (see Table 1) combining the two Gartner papers. Worth noting on the Gartner papers is that their data on systems administration differs between these papers by a significant amount: 1:106.7 versus 1:247.8. The papers did not explain the reasons for this difference.

As Table 1 shows, the help-desk area receives the most research. Almost every source of data has recommendations for help-desk staffing, and the mathematical mean of their recommendations is 1:87.3 (one help-desk staff for every 87.3 users) with a standard deviation of 20.5 users. Therefore, depending on the CNI's environment, the typical number of help-desk staff can range from 1:66.8 up through 1:107.8.

With three data points, systems administration is the next most heavily discussed area among the sources. The mathematical mean of the ratios is 1:169.9 with a standard deviation of 71.7 users. One advantage of the technical paper data is that companies often consider it more accurate than other sources of data. Therefore, when other sources of

	COTS Model Run A	Lucent Paper	Gartner Paper #1	Gartner Paper #2	DoD Org. Survey	DoD Expert Survey	Private Sector Expert Survey	COTS Model Run B	Mean	Std. Dev.
Users per FTE	51.6	50.2	31.6	48.2	33.0	43.8	44.9	38.5	42.1	7.3

Table 2: *Summary of Composite Ratios*

data start to show similarities to the technical papers, they tend to confirm each other's validity.

However, relative to other areas, the technical papers ignore both hardware maintenance (HM) and configuration management (CM). Only two papers contained hardware maintenance recommendations, and only one paper contained configuration management recommendations.

“If an organization's CNI staffing levels are too high, then it wastes resources. If the staffing levels are too low, then response times, reliability, and end-product quality suffers; overtime is too high; and workers leave for a better working environment.”

Organizational Survey

The organizational surveys represent data from existing organizations – some from the private sector and one from the Department of Defense (DoD). Unfortunately, in soliciting participation, none of the private-sector organizations were willing to participate openly. So MITRE submitted the survey anonymously to a different set of private-sector organizations in order to gain some unofficial information. MITRE obtained a few responses, but only two of them had enough clients and servers to be useful for this study. In general, the private-sector data has limited application since MITRE obtained data from small CNIs and obtained only two somewhat useful

responses. As for the DoD, MITRE ran into the same problem with the exception of one very large DoD organization, which was willing to share its information. Since this study focuses on private-sector data, one data point here was sufficient.

Table 1 also summarizes the data MITRE collected from its organizational surveys. The mean systems administration (SA) ratio (1:84.8) and the standard deviation (16.7 users) are much larger (i.e., more FTEs) than are those of the technical papers. The data seems to show a large disconnect between the technical papers and actual practice for SA. MITRE did not investigate potential causes of this difference, but one possibility is that the research centers are overly optimistic. Another possibility is that this data does not accurately reflect what organizations (in general) are actually doing (i.e., since the sample space is so small, it is not accurately showing the state of the practice). The numbers for help desk (HD) need no comment, since they are in general agreement. As for HM and CM, the ratios, again, are larger than those of the technical papers. Again, MITRE did not investigate the reasons for this difference but the same possibilities exist.

From Table 1, one can see how closely the ratios for SA, HD, and HM are to each other for each of the organizations. The private sector explains that they view the three areas as having overlapping talent because data points are such small CNIs. For the DoD data set, the organization has such specialized systems that they require a large number of SAs, thus pushing the SA ratio close to the other two ratios. Without some compelling need (such as the previous examples), an organization would not have as many SA staff as HD staff.

Technical-Expert Surveys

The technical-expert surveys represented best guesses at how experts might staff a sample CNI. For this survey, MITRE used a CNI containing approximately

Type of Data	Source of Data	Number of Users Per FTE of CNI Support			
		Systems Administration	Help Desk	Hardware Maintenance	Configuration Management
Technical Papers	1. Gartner #1 [4]	106.7	77.6	106.7	
	2. Gartner #2 [5-6]	247.8	60.0		
	3. PC Week [7]		86.0		
	4. IDC [8]		99.0		
Org. Surveys	6. DoD	103.3	110.7	106.9	442.9
	7. Private A	80.0	80.0	80.0	
	8. Private B	71.0		71.0	
Technical Expert Surveys	10. Private Sector	227.6	136.6	162.6	227.6
COTS Modeling Tools	12. Run B	80.8	199.4	193.1	292.8
Mean (Average)		131.0	106.2	120.0	321.1
Standard Deviation		74.2	44.3	48.0	110.4
Percent Standard Deviation		56.6	41.7	40.0	34.4

Table 3: *The Data Without Outliers*

100 servers, 1,000 clients, and 1,100 users. One expert from the private sector and one from the DoD answered the survey. Per the agreement on the survey, both respondents remain anonymous.

Table 1 (see page 22) summarizes the data from the two expert-opinion responses. The largest deviation between the two sets of answers is in the SA area, where the DoD expert's estimate is almost twice the private-sector experts estimate. There are also some drastic differences between the intra-organizational ratios. For example, the DoD response shows a clear spread between all four areas of CNI support with SA at the top of the graph and HD at the bottom. While the private-sector response is not as drastic, there is still a larger spread than found in the organizational survey.

Modeling

The modeling results represent data collected by taking the same scenario as the technical-expert surveys and running it through one of the well-known COTS modeling tools. Again, tool and vendor are anonymous.

Table 1 contains the data from two separate runs of the model: Run A views the CNI from a better light than run B (more details on this below). SA is very close between the runs. However, the other areas have a much larger deviation as the table shows. These differences are due to how the two sets of inputs characterized the scenario's CNI with respect to best practices and complexity, which are essential input parameters to the COTS model in question. Run A characterized

the scenario's CNI as more advanced with respect to best practices than run B. Run A also characterized the scenario's CNI as less complex than run B. The two runs provide some insight into how these parameters affect the modeling tool and thus affect the staffing levels, which the model predicts.

Note that the tool's values for HD are significantly different from all other sources of values for HD. With a HD ratio of 1:376, the author believes the tool is modeling more of a customer-serv-

“... the help-desk area receives the most research ... and the mathematical mean of their recommendations is 1:87.3 ... with a standard deviation of 20.5 users.”

ice center rather than a true help desk. The same holds true for HM. Since the vendor has no official tool validation, this issue may be a software error. Although this research strived to eliminate any differences between definitions, there may be a disconnect between the tool's terminology and those used in this article. That is, what the tool considers part of the HD support, this article may consider to be in some other category. These mapping issues are always a source of potential differences. However, these dif-

ferences go away when combining the data as Table 3 shows.

The Outliers

As the previous section mentioned, there appears to be some obvious outliers in the data. This section removes some of those outliers merely to show the effects on the data – both numerically and visually – since some organizations find such information useful.

Table 3 removes three outliers from the data set. The table drops the data from the Lucent paper due to its very small ratios for HM (1:284.5) and CM (1:853.5). The ratio of one CM person per 853 users is significantly different from all other data points for this area. Next, the table drops the DoD experts data, since it had a very small SA ratio (1:426.8) relative to all other data points. Lastly, the table drops the values from run A of the COTS modeling tool since it had very low ratios for HD, HM, and CM relative to the other sources.

Composite View of the Data

To remove potential differences between how the sources used terms (such as SA, HD, HM, and CM) and to provide an easy metric for predicting staffing sizes for CNI support, this section combines the data producing an overall FTE-to-user ratio. The author picked users (versus something like servers or clients), because most research in the field uses this same unit of measure (i.e., FTEs per number of users). In some cases, such as systems administration, logic dictates that a different unit of measure (e.g., FTEs per number of servers) is best; however, since the common unit of measure is users, the composite figures use it. In addition, all previous values use this unit of measure as well.

Table 2 contains the composite figures, but only for those sources of data that contained heuristics in more than one of the four areas. For example, the IDC paper as well as the *PC Week* paper referenced only help-desk staffing, so using these figures in a composite chart are not appropriate or useful. This table also ignores the two private-sector organizational surveys due to their small CNI size. The remaining eight sources of data provide CNI staffing ratios with a mean of 1:42.1, and whose standard deviation is just 7.3 – significantly better than the deviations from the non-composite ratios.

Issues

Before concluding, the reader should be aware of the issues MITRE encountered that affected the research. Some of these issues listed in the accompanying sidebar are specific to this particular study, but most of them are generic issues associated with labor studies involving CNI support. Despite these issues, however, the author believes that the resulting data, as a whole, gives a realistic and accurate picture of CNI staffing levels – both generically and specifically – since there is general agreement among the four areas of source data. That is, the technical papers (while having some outliers) are in agreement with the organizational data, the technical experts, and the modeling data from the COTS tool. Since writing this article, another DoD organization (of about 500 technical employees) has commented that its CNI staffing is approximately 1:42, which further supports the conclusions above. Another supporting factor is that after removing the obvious outliers, the deviations tighten up significantly. Nevertheless, some caution is appropriate since there is always a chance that the data just happen to agree.

Lastly, the data may not reflect optimal staffing levels: There is no oracle to tell us the optimum. Therefore, even if the data accurately portrays the state of the practice, it may not portray the optimum, since the state of the practice may not be optimal.

Conclusions

The method for determining support levels for CNIs is still an art, not a science as many would like. The lack of public information is partly due to the proprietary nature of company information. Despite the problems with collecting CNI staffing information, the data set as a whole appears accurate and useful, since there is general agreement among the four sources of data. That is, the technical papers, despite having some outliers, are in agreement with the organizational data, the experts' opinions, and the modeling data from the COTS tool. Another supporting factor is that after removing the obvious outliers, the deviations tighten up significantly, and the composite data is very tight for the newness of the industry.

Nevertheless, some caution is appropriate since there is always a chance that the data just happen to agree. Additionally, the data may not reflect

Research Issues to Consider

Following are some of the issues MITRE encountered that affected the research. Some of these issues are specific to this particular study, but most are generic issues associated with labor studies involving computer-networking infrastructure (CNI) support.

Time Constraints

MITRE limited this study to a three-month effort: January 2000 through March 2000. Within this timeline, there were further tradeoffs dealing with how much time to spend in each area of interest vs. the thoroughness of the analysis. For example, how much time to spend on technical papers vs. organizational surveys vs. expert-opinion surveys vs. modeling.

Technical Articles

While staffing levels are important organizational concerns, there are surprisingly very few technical papers on the subject. Of the technical papers that do exist, only a small number (six according to this study) discuss algorithms for determining FTEs with regard to CNI staffing. Of the six papers that do discuss algorithms, most focus only on a subset of the four major areas of CNI staffing.

Organizational Surveys

Another surprising outcome from this study was that very few companies were willing to share their CNI data.

Modeling Tools

A few companies claim to have modeling tools for calculating CNI staffing levels. The costs for the presumably better COTS modeling tools are quite high; therefore, MITRE used only one of the leading COTS tools in its study. Unfortunately, as the author found out during the research, the company who developed this COTS tool did not independently validate it, therefore, providing no confidence that it computed reasonable or accurate results. Also, some of the model's inputs, which should be essential factors in determining FTEs, are for "informational use only" according to the tool's manufacturer².

Mapping Data

Some of the data from the technical articles and the COTS tool required normalization to ensure that the data were in agreement (i.e., that the research counted apples as apples and oranges as oranges). Everyone seems to have slightly different definitions for the four primary areas of CNI support, which makes studying this area extremely difficult. Early on, MITRE learned that trying to make the areas of study too fine would prevent certain people from wanting to participate and would take too much time; therefore, MITRE kept the granularity at a high level (i.e., simple).

optimal staffing levels, and there is no oracle to tell us the optimum. Therefore, even if the data accurately portrays the state of the practice, it may not portray the optimum, since the state of the practice may not be optimal.

All charts in this report focus on user-based ratios for determining support levels (FTEs). However, there are other ratios: for example, those based on the number of servers and clients. When using these ratios, therefore, one must

ensure an accurate census before trying to estimate staffing levels. If one uses the user-based ratios, then that person or group must ensure an accurate accounting of users in the targeted organization beforehand.

Lastly, while these research findings focus on the private sector, they have application to any CNI. The most applicable ratio is the average overall FTE ratio of 1:42; that is, one FTE of CNI support for every 42 users with a standard

deviation of seven users. For example, one environment might have a ratio of around 1:35 (i.e., more support staff), while another environment would be 1:49 (fewer support staff). The deviation is about plus-or-minus 17.3 percent of the mean ratio. The HD ratios should also have close applicability to other domains, since the HD area received a lot of attention in the literature and seems to have strong agreement within both the literature and the surveys.

Of the remaining three CNI support areas – HM, SA, and CM – both the HM and SA ratios should provide rough estimates to other domains, while other domains may have trouble using the CM ratio. The state of the practice is very unclear with respect to CM, which is why applying the recommended CM ratio may be difficult and inaccurate for other domains. The state of the practice for HM and SA is more thorough but still not as solid as HD. Therefore, when applying HM and SA, other domains may need to allow for a wider variance than they would for HD.

The author hopes this article will be helpful to many DoD and non-DoD organizations trying to wrestle with this difficult and costly problem. Hopefully other organizations, because of the difficulties MITRE encountered, will share information more freely in the future. Lastly, the author encourages colleagues in the DoD and private sector to pass along any CNI staffing data whenever and wherever possible. While MITRE collected all of the technical articles they could find, the author would appreciate hearing about any significant references that our searches may have missed, i.e.,

anything not listed in the references section. ♦

References

1. Weinberg, Neal. "TCO Tall Tales." *Network World* 7 June 1999.
2. Gowen, Lon. "Computer-Networking Infrastructure Manpower Study." The MITRE Corporation, Document MP-00B0000018. Apr. 2000 (not in the public domain).
3. "Network Operations Center Staffing Plan." Lucent International Network Services, Feb. 1999.
4. Silver, M. "The Staffing Connection to PC/LAN TCO." Gartner Group. 1 May 1998.
5. Cappuccio, D. "The Metrics of LAN Staffing – A Top-Down View." Gartner Group. Nov. 1996.
6. Apfel, A., K. McGee, J. Pultz, A. Schoeller, M. Zboray, and Ardito C. Smith, eds. "Selection Criteria for Tomorrow's Enterprise Networking Professionals." Gartner Group. 8 July 1998.
7. Plotnick, Neil. "Getting Your People Power in Perspective." *PC Week*, 12 Apr. 1999.
8. Kavanagh, Kelly, Tom Oleson, and Chris Hoffman, eds. "A Model for Determining Help-Desk Staffing." IDC Government. July 1998.

Notes

1. TCO is a term for which there is no "accepted industry standard"; however, the term usually includes just what its name says – all costs associated with owning a piece of hardware, including the support and maintenance.

2. An excellent area for research, therefore, would be 1) to compare as many of these models against each other as possible and 2) to determine their accuracy (i.e., attempt some sort of validation). Currently, there are no analyses in the literature (that the author could find) for any of these models.

About the Author




Lon D. Gowen, Ph.D., is a lead staff engineer for The MITRE Corporation. He currently works on a MITRE contract at the U.S. Strategic Command, Offutt AFB, Neb.

Gowen has doctorate and master's degrees from Arizona State University with areas of interest in software engineering, software languages, and embedded systems. He has a bachelor's degree with a double major in computer science and applied mathematics from the University of Nebraska. Gowen was a Distinguished Visitor for the Institute of Electrical and Electronics Engineers Computer Society for three years. He also teaches the systems and software safety classes for the NASA Safety Training Center.

The MITRE Corporation
1004 Lincoln Road, PMB 327
Bellevue, NE 68005
Phone: (402) 294-2689
Fax: (402) 294-1264
E-mail: gowenld@mitre.org

We accept article submissions on all software-related topics at any time.
Please follow the Author Guidelines for CROSSTALK, available on the Internet at:
www.stsc.hill.af.mil/crosstalk/xtlkguid.pdf



Call for Articles

If your experience or research has produced information that could be useful to others, **CROSSTALK** can get the word out. We are especially looking for articles in several specific, high-interest areas. Upcoming issues of **CROSSTALK** will have special, yet nonexclusive, focuses on the following tentative themes:

System Requirement Risks

March 2002

Submission Deadline: Oct. 24, 2001

Software Estimation

April 2002

Submission Deadline: Nov. 21, 2001

Information Assurance

June 2002

Submission Deadline: Jan. 23, 2002